

CLAIMS

What is claimed is:

1. A method for storing data on a plurality of disk drives comprising:
 - addressing a plurality of data strips from said data to a chosen disk of said plurality of disk drives;
 - forming a data stream comprising said data strips, said data stream having a first throughput;
 - creating a plurality of parallel data streams, each of said plurality of parallel data streams having a second throughput, said second throughput being smaller than said first throughput;
 - directing said plurality of parallel data streams to a corresponding plurality of said plurality of disk drives such that each data strip in said plurality of data strips is transmitted to said chosen disk of said plurality of disk drives; and
 - storing each of said data strips on said each of plurality of disk drives.
2. The method of claim 1 wherein the number of said plurality of parallel data streams is approximately equal to said first throughput divided by said second throughput.
3. The method of claim 2 wherein said number of said plurality of parallel data streams is 2.
4. The method of claim 2 wherein said number of said plurality of parallel data streams is 4.
5. The method of claim 1 wherein at least one of said data strips comprises parity information.
6. The method of claim 1 wherein said creating a plurality of parallel data streams is performed by a first-in-first-out buffer.
7. The method of claim 1 wherein said directing said plurality of parallel data streams is performed by a crossbar switch.
8. The method of claim 1 further comprising:
 - reading said each of said data strips from said plurality of disk drives;

transmitting said each of said data strips from said plurality of disk drives
in said plurality of parallel data streams; and
transmitting said first data stream containing said data strips.

9. A system for storing data comprising:

5 a plurality of disk drives each having a communication channel capable
of communicating at a first throughput;

a controller adapted to address a plurality of data strips from said data to
a chosen disk of said plurality of disk drives, and form a data stream
comprising said data strips, said data stream having a second throughput;

10 a buffered switch in communication with said controller and adapted to
create a plurality of parallel data streams, each of said plurality of parallel data
streams having a second throughput, said first throughput being smaller than
said second throughput;

15 a crossbar switch in communication with said buffered switch and
adapted to direct said plurality of parallel data streams to a corresponding
plurality of said plurality of disk drives such that each of said separate data
strips are transmitted to each of said plurality of disk drives to which said
separate data strips are addressed; and

20 wherein said plurality of disk drives are adapted to receive said plurality
of parallel data streams and store said data strips on said disk drives.

10. The system of claim 9 wherein the number of said plurality of parallel data
streams is approximately equal to said second throughput divided by said first
throughput.

25 11. The system of claim 10 wherein said number of said plurality of parallel data
streams is 2.

12. The system of claim 10 wherein said number of said plurality of parallel data
streams is 4.

13. The system of claim 9 wherein at least one of said data strips comprises
parity information.

30 14. A system for storing data comprising:

a plurality of storage means each having a communication channel capable of communicating at a first throughput;

a controlling means adapted to determine a first data stripe to store onto said plurality of disk drives, said data stripe containing a separate data strip addressed to each of said plurality of disk drives, and adapted to provide a first data stream having a second throughput and containing said data strips;

a buffer means in communication with said controlling means and adapted to receive said first data stream, said buffer means having a first-in-first-out buffer into which said first data stream is received at said second throughput, said buffer means further adapted to remove said data strips from said first-in-first-out buffer to create a plurality of parallel data streams, each of said plurality of parallel data streams having said first throughput, said first throughput being smaller than said first throughput;

a switch means in communication with said buffer means and adapted to direct said plurality of parallel data streams to a corresponding plurality of said plurality of disk drives such that each of said separate data strips are transmitted to each of said plurality of disk drives to which said separate data strips are addressed; and

wherein said storage means are adapted to receive said plurality of data streams and store said data strips on said storage means.

15. The system of claim 14 wherein the number of said plurality of parallel data streams is approximately equal to said second throughput divided by said first throughput.

16. The system of claim 15 wherein said number of said plurality of parallel data streams is 2.

17. The system of claim 15 wherein said number of said plurality of parallel data streams is 4.

18. The system of claim 14 wherein at least one of said data strips comprises parity information.